

Distributed Time-Varying Stochastic Optimization and Utility-based Communication

Andrea Simonetto, Leon Kester, and Geert Leus

Abstract—We devise a distributed asynchronous stochastic ϵ -gradient-based algorithm to enable a network of computing and communicating nodes to solve a constrained discrete-time time-varying stochastic convex optimization problem. Each node updates its own decision variable only once every discrete time step. Under some assumptions (among which, strong convexity, Lipschitz continuity of the gradient, persistent excitation), we prove the algorithm's asymptotic convergence in expectation to an error bound whose size is related to the constant stepsize choice α , the variability in time of the optimization problem, and to the accuracy ϵ . Moreover, the convergence rate is linear. Then, we show how to compute locally stochastic ϵ -gradients that depend also on the time-varying noise probability density function (PDF) of the neighboring nodes, without requiring the neighbors to send such PDFs at each time step. We devise utility-based policies to allow each node to decide whether to send or not the most up-to-date PDF, which guarantee a given user-specified error level ϵ in the computation of the stochastic ϵ -gradient. Numerical simulations display the added value of the proposed approach and its relevance for estimation and control of time-varying processes and networked systems.

I. INTRODUCTION

We consider a time-varying stochastic optimization problem defined on time-varying functions that are distributed over a network of computing and communicating nodes. Let the nodes be labeled with $i \in V = \{1, \dots, n\}$, and for each discrete time $k \in \mathbb{N}$, we equip each of them with the local function $f_{i,k}(\mathbf{x}, \boldsymbol{\omega}) : \mathbb{R}^{d \times q} \rightarrow \mathbb{R}$. In particular, the common vector $\mathbf{x} \in \mathbb{R}^d$ represents the decision variables, while the stochastic vector $\boldsymbol{\omega} \in \Omega \subseteq \mathbb{R}^q$ is a stochastic variable drawn from a given (or estimated) time-varying probability density function (PDF) $p_{\omega,k}(\boldsymbol{\omega})$ on Ω . We assume that $\boldsymbol{\omega}$ is comprised of local stochastic variables as $\boldsymbol{\omega} = (\omega_1^T, \dots, \omega_n^T)^T$, each of them of possible different dimensions and each of them uncorrelated with one another. Furthermore, we let $\omega_i \in \Omega_i$, $\Omega = \prod_{i=1}^n \Omega_i$, and $\omega_i \sim p_{\omega_i,k}(\omega_i)$ on Ω_i .

The main goal for the computing nodes at each discrete time k is to solve the optimization problem

$$\underset{\mathbf{x} \in X_k}{\text{minimize}} \mathbb{E}_{\omega,k} \left[\sum_{i \in V} f_{i,k}(\mathbf{x}, \boldsymbol{\omega}) \right] := \int_{\Omega} \sum_{i \in V} f_{i,k}(\mathbf{x}, \boldsymbol{\omega}) p_{\omega,k}(\boldsymbol{\omega}) d\boldsymbol{\omega}, \quad (1)$$

where each of the $f_{i,k}(\mathbf{x}, \boldsymbol{\omega})$ is a convex function of \mathbf{x} for all $\boldsymbol{\omega} \in \Omega$, while X_k is a compact convex set.

We allow the computing nodes to communicate with their immediate neighbors defined via the undirected communication graph $G = (V, E)$, with edge set E . In particular, each

node i can communicate with all the nodes $j \in N_i := \{j \in V | (i, j) \in E\}$. This communication possibility is not assumed to be synchronous among the computing nodes, but can happen at asynchronous times. We only assume that when node i communicates with node j , node j communicates with node i as well (that is, we assume an edge-asynchronous protocol).

Stochastic optimization problems like (1) are rather popular in machine learning [1], [2]. In this paper, we are more interested in their connection to the field of reasoning under uncertainty, self-awareness, and optimization-based design for distributed systems [3], [4], where each computing node may learn or change the estimate of its local $p_{\omega_i,k}$ with time. In addition to that, other instances of the problem encompass distributed stochastic control [5], finance [6], and distributed statistical signal processing [7]. As for possible high-impact societal applications, distributed time-varying stochastic problems as (1) arise in the context of smart grids [8], [9] and cooperative adaptive cruise control [10], to name a few.

If, for the moment, we focus only at the time-varying nature of the problem, we can also recognize in (1) optimization programs that appear in distributed estimation of stochastic time-varying signals [7], in distributed control of mobile multi-robot systems with time-varying tasks [11], and as a result of sequential convex programming approaches to multi-agent non-convex problems [12]. When each of the functions $f_{i,k}(\mathbf{x})$ and the set X_k are time-invariant, several approaches can be applied to solve (1). These techniques differ for the assumptions they require and the properties they can ensure (convergence, convergence rate, resilience to asynchronous communication protocols, among others). Examples of such approaches are the stochastic subgradient [13], dual averaging [14], and the alternating direction method of multipliers [15]. Since the aforementioned techniques are iterative, and they require communication among the nodes to converge to an optimizer of (1), they would *provably* converge in the case of time-varying $f_{i,k}(\mathbf{x})$'s and/or time-varying X_k only when each node could exchange an infinite number of messages with its neighbors, between consecutive time steps k and $k + 1$. Specific methods that account for a finite number of messages between consecutive time steps and still guarantee convergence have been proposed in [7], [11], [16]–[23], but they are all limited to specific *deterministic* versions of (1). Notably, in [7], [23], the authors work under the same assumptions that we will use, however they consider deterministic and *unconstrained* optimization problems, while in [22], the authors employ deterministic subgradient methods and assume that the optimizers of (1) do not change in time.

Contributions. As a first contribution, we consider the stochastic nature of the optimization problem and we propose an asynchronous stochastic gradient-based distributed algo-

A. Simonetto and G. Leus are with the Faculty of EEMCS, Delft University of Technology, 2826CD Delft, The Netherlands. e-mails: {a.simonetto, g.j.t.leus}@tudelft.nl. L. Kester is with TNO, Oude Waalsdoorperweg 63, 2597AK The Hague, The Netherlands. email: leon.kester@tno.nl. This research was supported in part by STW under the D2S2 project from the ASSYS program (project 10561).

rithm for the computing nodes to converge to an optimizer of (1). In fact, due to the time-varying nature of the problem and due to possible errors in the computation of the stochastic gradient, the convergence will be shown up to an error bound, whose size is directly dependent on the mentioned elements. This algorithm can be seen as a generalization in a time-varying context of the work in [13] where only one iteration of the algorithm is performed between consecutive time steps, as well as a generalization of the work in [7], [23] in constrained, asynchronous, and stochastic settings. In addition, in contrast to [7], [23], our algorithm does not hinge on dual variables to reach a common decision vector among the nodes (which complicates significantly the theoretical analysis of convergence), but is instead based on consensus protocols, which are easier to analyze and embed on real hardware.

The proposed algorithm can also be seen as a stand-alone contribution, that can be applied to many different scenarios, for example, in deterministic settings or when the nodes can compute the gradient exactly (we will show an example of these cases in the numerical simulations).

As a second contribution, we consider the case in which the computing nodes have access to the local PDFs $p_{\omega_i,k}$ only. In this case, to avoid the communication overhead of sending at each time step k , the most current local PDF to the other nodes, we devise a utility-based policy that enables each node to decide locally whether to send such $p_{\omega_i,k}$ or not. This utility-based policy guarantees that despite the possible outdated information on $p_{\omega_i,k}$, each of the nodes is able to compute their stochastic gradient up to a user-specified precision ϵ . This fits perfectly in the scheme of the proposed asynchronous distributed algorithm, and is particularly useful to limit the overhead of sending at each time step the varying PDFs, which might be non-Gaussian and difficult to encode in a message with a limited number of parameters. This second contribution can be seen as a generalization of event-triggered optimization [24], where we determine the triggering mechanism not only based on convergence arguments, but also on performance guarantees.

Interesting related work for event-based mechanisms can be found in event-triggered control [25], [26], estimation [27], and energy constrained communication networks [28], [29]. An important difference between our contribution and these fields of research is that we cast the global optimization problem as a stochastic program, which is able to capture *unpredicted (i.e., non-modeled) but locally measurable changes* and judge *locally* their importance to the *global* performance.

Organization. The first contribution of the paper is encoded in Problem 1, which is formally expressed in Section III and whose solution is outlined in Algorithm 1 and characterized in Theorem 1. The second contribution is the approach to tackle Problem 2, which is fully discussed in Section IV and Theorem 2. Numerical results support the theoretical findings and are reported in Section V. All the proofs are grouped in Section VI.

Notation. For any vector $\mathbf{x} \in \mathbb{R}^n$, the norm $\|\mathbf{x}\|$ represents the standard Euclidean norm. For any real-valued squared matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$, we say $\mathbf{X} \geq 0$ or $\mathbf{X} \leq 0$ if the matrix is positive semi-definite or negative semi-definite, respectively. For any positive semi-definite matrix of dimension n , we order

its eigenvalues as $0 \leq \lambda_1(\mathbf{X}) \leq \dots \leq \lambda_n(\mathbf{X})$. The symbol \otimes represents the Kronecker product. Given a differentiable convex function $f(\mathbf{x}) : X \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$ and a non-negative scalar ϵ , an ϵ -(sub)gradient of $f(\mathbf{x})$ at $\mathbf{x} \in X$ is a vector $\tilde{\mathbf{g}} \in \mathbb{R}^n$ such that

$$\tilde{\mathbf{g}}^\top (\mathbf{y} - \mathbf{x}) \leq f(\mathbf{y}) - f(\mathbf{x}) + \epsilon, \quad \text{for all } \mathbf{y} \in X, \quad (2)$$

and finally, given a compact set X , the symbol $|X|$ stays for $|X| := \max_{\mathbf{x} \in X} \{\|\mathbf{x}\|\}$.

II. ASSUMPTIONS AND PROBLEM STATEMENT

Let $X_{k,k-1}$ be defined as the convex hull of the union of the sets X_k and X_{k-1} , for $k \geq 1$, i.e., $X_{k,k-1} := \text{cvxh}\{X_k \cup X_{k-1}\}$, which will be used in our proofs. Throughout the text we assume the following simplifying assumptions.

Assumption 1: (Sets and Objective functions) The sets X_k are compact convex sets for all $k \geq 0$. The functions $f_{i,k}(\mathbf{x}, \omega)$, $k \geq 1$, are twice differentiable and convex for all $\mathbf{x} \in X_{k,k-1}$, $\omega \in \Omega$, and they have the following properties:

- 1) the expectations $\mathbb{E}_{\omega,k}[f_{i,k}(\mathbf{x}, \omega)]$ are strongly convex with respect to $\mathbf{x} \in X_{k,k-1}$, $k \geq 1$. In particular, the eigenvalues of the Hessian $\mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}\mathbf{x}}^2 f_{i,k}(\mathbf{x}, \omega)]$ are lower bounded by the strong convexity constant m_f for all $\mathbf{x} \in X_{k,k-1}$,

$$\mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}\mathbf{x}}^2 f_{i,k}(\mathbf{x}, \omega)] \geq m_f \mathbf{I},$$

or equivalently

$$\begin{aligned} \mathbb{E}_{\omega,k}[f_{i,k}(\mathbf{y}, \omega) - f_{i,k}(\mathbf{x}, \omega)] &\geq \\ \mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega)^\top](\mathbf{y} - \mathbf{x}) + \frac{m_f}{2} \|\mathbf{y} - \mathbf{x}\|^2, \end{aligned}$$

for all vectors \mathbf{x}, \mathbf{y} in $X_{k,k-1}$;

- 2) the gradients of the expectations $\mathbb{E}_{\omega,k}[f_{i,k}(\mathbf{x}, \omega)]$ with respect to \mathbf{x} are Lipschitz continuous with constant L , meaning that

$$\|\mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega) - \nabla_{\mathbf{y}} f_{i,k}(\mathbf{y}, \omega)]\| \leq L \|\mathbf{x} - \mathbf{y}\|,$$

for all vectors \mathbf{x}, \mathbf{y} in $X_{k,k-1}$, or equivalently

$$\mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}\mathbf{x}}^2 f_{i,k}(\mathbf{x}, \omega)] \leq L \mathbf{I},$$

for all $\mathbf{x} \in X_{k,k-1}$.

Assumption 2: (Stochastic variables) The stochastic variables ω_i affect the objective functions only locally, i.e.,

$$\nabla_{\omega_j} f_{i,k}(\mathbf{x}, \omega) = \mathbf{0}, \quad \text{if } j \notin N_i \cup i;$$

the local cost functions are separable w.r.t. the ω_i , i.e.,

$$f_{i,k}(\mathbf{x}, \omega) = \sum_{j \in N_i \cup i} f_{i,k,j}(\mathbf{x}, \omega_j), \quad \text{for } i = 1, \dots, n;$$

and finally, each of the PDFs $p_{\omega,k}$ are defined over a compact set, that is to say that each of the Ω_i is a compact set.

Assumption 3: (Optimal problem) There exists a solution \mathbf{x}_k^* for (1) at each time step k , and the distance between these solutions at two subsequent time steps is upper bounded as

$$\|\mathbf{x}_{k+1}^* - \mathbf{x}_k^*\| \leq \delta_{\mathbf{x}}, \quad \text{for } k \geq 1,$$

where $\delta_{\mathbf{x}}$ is a finite non-negative scalar.

Assumption 4: (Communication graph) At each iteration k the symmetric adjacency matrix of the communication graph \mathbf{A}_k is generated by an i.i.d. Bernoulli process with $\Pr[\mathbf{A}_k]_{ij} = 1] = s_{ij} > 0$ for all $(i, j) \in E$, with a given probability $0 < s_{ij} \leq 1$. Furthermore, let E_k be the communication edge set at time k , and let $G_k := (V, E_k)$. For every $k' \geq 1$, there exists an integer $T \geq 0$ such that:

- 1) the union of the edge sets satisfies $\bigcup_{\ell=k'}^{k'+T} E_\ell = E$;
- 2) the union graph, i.e., $\bigcup_{\ell=k'}^{k'+T} G_\ell$, is connected.

Assumptions 1 and 3 are typical in time-varying optimization problems; in particular, Assumption 1 ensures that the optimizer of (1) (if it exists) is unique, while Assumption 3 ensures the existence of said optimizers and bounds their variability, which will be crucial in the convergence of the proposed solution. For a broader perspective on these two assumptions, the reader is referred to [30, Chapter 6]. We notice the need of considering a larger set $(X_k \cup X_{k-1})$ for the validity of the smoothness conditions on the cost functions; this is due to the time-varying nature of the problem and its role will be clear in the proofs. The compactness assumption on X_k is also quite used in (ϵ) -subgradient algorithms and it is not very restrictive (one could always use a box constraint to limit the variability of the decision variable \mathbf{x}); furthermore, this assumption is not necessary when the computation of the gradient is exact. Assumption 4 is typically required in asynchronous distributed protocols and is rather weak. Assumption 2 is a simplifying assumption on the nature of the stochastic interdependences; in particular Assumption 2 says that only the neighboring stochastic variables have an effect on each node's cost function $f_{i,k}$, and that this effect is the sum of local components coming from different neighbors. This is certainly quite reasonable in distributed systems. Finally, Assumption 2 says that the probability space is compact, which is not a very restrictive requirement in practice and it is important in quantifying how important are the variations in $f_{i,k}$ w.r.t. the variations in $p_{\omega,k}$.

As mentioned in the Introduction, the problem we are interested in is twofold. First of all, we want to enable the computing nodes to solve (1) in a distributed fashion, where each of the nodes communicates with their neighbors only. For this first task, we introduce local copies of the decision variable \mathbf{x}_k . These local copies are referred to as $\mathbf{y}_{i,k}$. We formally formulate the problem at hand as

Problem 1: Devising a stochastic ϵ -(sub)gradient distributed algorithm in order to enforce that the local decision variable $\mathbf{y}_{i,k}$ eventually converges (up to a bounded error) to the optimal solution of (1) at time step k (\mathbf{x}_k^*), or formally,

$$\liminf_{k \rightarrow \infty} \mathbb{E} [\|\mathbf{y}_{i,k} - \mathbf{x}_k^*\|^2] \leq \delta, \text{ for } i \in V,$$

for some $\delta \geq 0$, which has to decrease if $\delta_x \rightarrow 0$ and $\epsilon \rightarrow 0$.

The second aspect of the problem we consider is to limit the need of communicating the most actual PDF $p_{\omega,k}$ among the nodes and can be formulated as

Problem 2: Devising a utility-based policy that allows the nodes to decide whether to send or not their $p_{\omega_i,k}$ and yet that guarantees a prescribed level of accuracy ϵ in the computation of the stochastic ϵ -(sub)gradient.

III. PROPOSED APPROACH FOR PROBLEM 1

In order to solve the first part of the problem we propose a distributed asynchronous stochastic gradient algorithm. We start by defining the time-varying matrix \mathbf{W}_k and two different stepsizes $\alpha > 0$ and $\beta > 0$. The matrix \mathbf{W}_k is a symmetric matrix based on the adjacency matrix \mathbf{A}_k as,

$$[\mathbf{W}_k]_{i,j} = \begin{cases} -[\mathbf{A}_k]_{i,j} & \text{for } j \neq i \\ \sum_{l=1}^n [\mathbf{A}_k]_{i,l} & \text{for } j = i \end{cases}. \quad (3)$$

From Assumption 4, the symmetric matrix \mathbf{W}_k has nonzero elements if and only if the related nodes can communicate with each other, it is rank deficient and in particular $\mathbf{W}_k \mathbf{1}_n = \mathbf{0}_n$, and finally, for the sequence of matrices $\{\mathbf{W}_k\}$,

$$\mathbb{E}[\mathbf{W}_k] = \bar{\mathbf{W}} = \bar{\mathbf{W}}^T, \text{ with } \lambda_2(\bar{\mathbf{W}}) > 0. \quad (4)$$

As for the stepsizes α and β , the former caters for convergence to the optimal solution of (1), while the latter dictates the consensus among the different computing nodes. The proposed algorithm is outlined in Algorithm 1.

Algorithm 1 Asynchronous distributed stochastic ϵ -gradient

Initialize by picking locally an arbitrary $\mathbf{y}_{i,1} \in X_0$. Then for $k \geq 1$:

- 1) compute the local variable $\mathbf{v}_{i,k+1}$ by local communication as

$$\mathbf{v}_{i,k+1} = \mathbf{y}_{i,k} - \beta \sum_{j=1}^n [\mathbf{W}_k]_{i,j} \mathbf{y}_{j,k}, \quad (5a)$$

for which we will show that in fact $\mathbf{v}_{i,k+1} \in X_{k-1}$ in Lemma 1;

- 2) compute locally the stochastic ϵ -(sub)gradient of $f_{i,k}$ with respect to \mathbf{x} at $\mathbf{v}_{i,k+1}$, as

$$\begin{aligned} \tilde{\mathbf{g}}_{i,k} &= \mathbb{E}_{\omega'} [\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega) | \mathbf{v}_{i,k+1}] := \\ &\int_{\Omega} (\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega) | \mathbf{v}_{i,k+1}) p_{\omega'}(\omega) d\omega, \end{aligned} \quad (5b)$$

for which $p_{\omega'} \in \Omega$ is the PDF of the stochastic variable ω known by node i at time k (which is possibly outdated). We give a sufficient condition for the chosen $\tilde{\mathbf{g}}_{i,k}$ to be a stochastic ϵ -(sub)gradient in Lemma 1;

- 3) update the local variable $\mathbf{y}_{i,k}$ as

$$\mathbf{y}_{i,k+1} = \mathbf{P}_{X_k} [\mathbf{v}_{i,k+1} - \alpha \tilde{\mathbf{g}}_{i,k}]; \quad (5c)$$

where $\mathbf{P}_{X_k}[\cdot]$ indicates the projection operator;

- 4) go to step 1.
-

The proposed algorithm involves only local communication (in step 1) and local computations. This is true, provided that each node can obtain the PDF p_{ω}' necessary in step 2. We show later in this section how this can be achieved by a utility-based local communication mechanism, that can be completely disconnected by the communication of the local $\mathbf{y}_{i,k}$'s.

Our first result establishes convergence of the sequence of local variables $\{\mathbf{y}_{i,k}\}$ generated by Algorithm 1 to the optimizer of (1) up to a bounded error. The error is comprised of two parts, one coming from the error in the computation of the gradient, the other coming from the time-varying nature of the original problem.

As for Assumptions 1 and 3, the gradient of the expectations at optimality are bounded for all k , and in particular,

$$\max_{i \in V} \left\{ \left\| \sum_{j \in V, j \neq i} \mathbb{E}_{\omega,k} [\nabla_{\mathbf{x}} f_{j,k}(\mathbf{x}, \omega) | \mathbf{x}_k^*] \right\| \right\} \leq G, \quad (6)$$

for a certain nonnegative bound G .

Theorem 1: Let $\{\mathbf{y}_{i,k}\}$ be the sequence generated by Algorithm 1 where each $\tilde{\mathbf{g}}_{i,k}$ is a stochastic ϵ -(sub)gradient of $f_{i,k}$ w.r.t. \mathbf{x} at $\mathbf{v}_{i,k+1}$ computed as (5b); let $\mathbf{g}_{i,k}$ be the up-to-date stochastic gradient as defined in (15); let \mathbf{y}_k be the stacked nd -dimensional column vector of the local variables $\mathbf{y}_{i,k}$. Let G be defined as (6). Define $\varrho := 1 + \alpha^2 L^2 - \alpha m_f$. Under Assumptions 1 till 4, if $\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \epsilon/2|X_{k,k-1}|$ for each $k \geq 1$, then by choosing $\beta < 1/n$, and $\alpha < m_f/L^2$, the sequence $\{\mathbf{y}_k\}$ converges as

$$\liminf_{k \rightarrow \infty} \mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{1}_n \otimes \mathbf{x}_{k+1}^*\|^2] \leq \frac{1}{1-\varrho} \left(\alpha \hat{\psi}(\alpha, \epsilon, \sqrt{\gamma}) \frac{1}{\sqrt{\gamma}} + n \delta_x^2 \frac{1}{1-\sqrt{\gamma}} \right),$$

where $\gamma = 1 - \beta \lambda_2(\bar{\mathbf{W}})$ and

$$\hat{\psi}(\alpha, \epsilon, \sqrt{\gamma}) = n \frac{\epsilon}{\sqrt{\gamma}} \left(\alpha \frac{\epsilon}{4|X_{k,k-1}|^2} + 2\alpha L + 2 \right) + \frac{\alpha n G^2}{1-\sqrt{\gamma}}.$$

Furthermore, $0 < \varrho < 1$ and the convergence rate is linear.

As we will see in Section VI, the proof of Theorem 1 is based on the Peter-Paul inequality for a particular choice of the parameter μ . This choice is valid if $0 < \gamma < 1$, as in our case, and allows us to separate the effect of α with the effect of the network connectivity $\lambda_2(\bar{\mathbf{W}})$. Furthermore, it gives us some intuition on which errors will play a more predominant role for different $\lambda_2(\bar{\mathbf{W}})$. In particular, for highly connected graphs, when $\gamma \rightarrow 0$, then the term $\hat{\psi}(\alpha, \epsilon, \sqrt{\gamma})$ is more important than the time-varying term. This is expected since consensus mixing is faster and the nodes quickly agree on a common approximate optimizer. Furthermore, in this case, to reduce the error, we will have to choose smaller and smaller α 's. On the contrary for poorly connected graphs, $\gamma \rightarrow 1$, the error term with G and the time-varying one are dominant. A similar results was found in [7] for deterministic strongly convex optimization problems using a *synchronized* dual decomposition approach.

Remark 1: The bound that has been derived is not expected to be tight for the whole range of γ 's; however one can use the same Peter-Paul inequality with a different choice of μ to obtain tighter bounds for specific values of γ . Our choice was determined based on the idea of dividing the terms of convergence (α) and consensus (β). In addition, it is expected that the smaller δ_x and ϵ are, the tighter the bound is.

IV. PROPOSED APPROACH FOR PROBLEM 2

We focus now on the second part of our problem, which is how to generate ϵ -(sub)gradient vectors for a specific value of ϵ , that is, under the sufficient condition of Lemma 1 in Section VI, how to generate $\tilde{\mathbf{g}}_{i,k}$ such that $\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \epsilon/2|X_{k,k-1}|$. In particular, we consider a utility-based event-triggered mechanism that allows the computing nodes to send one another the latest updated PDF of ω only when needed, to guarantee a prescribed level of ϵ .

The error in computing the stochastic gradient at node i is

$$\begin{aligned} \mathbf{g}_{i,k} - \tilde{\mathbf{g}}_{i,k} &= \mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega)|_{\mathbf{v}_{i,k+1}}] - \mathbb{E}_{\omega'}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega)|_{\mathbf{v}_{i,k+1}}] = \\ &= \sum_{j \in N_i} \mathbb{E}_{\omega_j,k}[\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \mathbb{E}_{\omega'_j}[\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] \end{aligned}$$




Fig. 1. Quantities known at each node and the communication policy π_{ij} , the quantities with ' are the outdated ones.

and as a consequence of Lemma 1, we can derive the sufficient condition for each $j \in N_i$

$$\begin{aligned} \|\mathbb{E}_{\omega_j,k}[\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \mathbb{E}_{\omega'_j}[\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}]\| \\ \leq \frac{\epsilon}{2|X_{k,k-1}|\Delta_i}, \end{aligned} \quad (7)$$

where Δ_i is the number of possible neighbors of i , as in E .

If we could use the condition (7) to establish a sending policy, then $\tilde{\mathbf{g}}_{i,k}$ would be guaranteed to be an ϵ -(sub)gradient and we would know the error floor in Theorem 1. Unfortunately, neither node i nor node j can compute, check, or enforce (7). This is because nodes have either their updated $\mathbf{v}_{i,k+1}$ or the updated $p_{\omega_{j,k}}$, but not both at the same time (see Figure 1 for a pictorial representation). However, along the mechanism to send the most up-to-date $p_{\omega_{j,k}}$, we can think of a similar mechanism to send or not the most up-to-date gradient $\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}$, which is a function of ω_j . Let $\nabla_{\mathbf{x}} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}$ be the possibly outdated gradient function of node i available at node j . We define the utility metrics,

$$\begin{aligned} \mathcal{U}_{S,ij}^{(1)}(\omega_k|\omega') &:= \|\mathbb{E}_{\omega_j,k}[\nabla_{\mathbf{x}} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}] - \mathbb{E}_{\omega'_j}[\nabla_{\mathbf{x}} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\|, \end{aligned} \quad (8a)$$

$$\mathcal{U}_{S,j}^{(2)}(\omega_k|\omega') := \max_{\omega_j \in \Omega_j} |p_{\omega_{j,k}}(\omega) - p_{\omega'_j}(\omega)|, \quad (8b)$$

$$\begin{aligned} \mathcal{U}_{R,ij}(\nabla_{\mathbf{x}} f_{i,k,j}|\nabla_{\mathbf{x}} f'_{i,j}) &:= \int_{\Omega_j} \|\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}} - \nabla_{\mathbf{x}} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}\| d\omega_j. \end{aligned} \quad (8c)$$

These utility functions specify how useful is the up-to-date information with respect to the outdated one. All of them are locally computable by nodes i and j , respectively (Figure 1). With these quantities, we specify the sending policy π_{ij} of node i for each neighboring node j , for any scalar $\eta \in [0, 1]$ and $\nu > 0$ as follows:

$$\pi_{ij} : \begin{cases} \text{iff } \mathcal{U}_{R,ij}(\nabla_{\mathbf{x}} f_{i,k,j}|\nabla_{\mathbf{x}} f'_{i,j}) > \frac{\eta}{\nu} \frac{\epsilon}{2|X_{k,k-1}|\Delta_i}, \\ \quad \text{send } \nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}} \text{ to } j \\ \text{iff } \left[\begin{array}{l} \text{either } \mathcal{U}_{S,ij}^{(1)}(\omega_k|\omega') > (1-\eta) \frac{\epsilon}{2|X_{k,k-1}|\Delta_j} \\ \text{or } \mathcal{U}_{S,i}^{(2)}(\omega_k|\omega') > \nu \end{array} \right], \\ \quad \text{send the PDF } p_{\omega_{i,k}} \text{ to } j. \end{cases} \quad (9)$$

With this in place, the second main result of the paper can be formalized.

Theorem 2: Let each computing node i use the policy π_{ij} (for each j , such that $(i, j) \in E$) defined in (9) for specified values of ϵ, η , and ν . Under Assumptions 1 till Assumption 4 and with the same definitions of Theorem 1, the quantity $\hat{\mathbf{g}}_{i,k}$ defined in (5b) satisfies $\|\hat{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \epsilon/2 |X_{k,k-1}|$ and it is a stochastic ϵ -(sub)gradient of $f_{i,k}(\mathbf{x}, \boldsymbol{\omega})$ w.r.t. $\mathbf{v}_{i,k+1}$.

Theorem 2 guarantees that our proposed utility-based policy will deliver ϵ -gradient vectors of a prescribed accuracy, thereby ensuring convergence in the sense of Theorem 1. The policy π_{ij} is completely disconnected with the running of Algorithm 1, and it can be triggered at any time step k . Its utility-based nature allows each node to send its neighbors the up-to-date information only when they really need it. We remark that the information that each node has to send is the gradient of $f_{i,k,j}$ with respect to \mathbf{x} at $\mathbf{v}_{i,k+1}$, and the PDF of $\boldsymbol{\omega}_{i,k}$. Both are functions of $\boldsymbol{\omega}_{i,k}$, which need to be encoded in some way. We will show with a simple numerical evaluation how this can be done in practice and that the added communication cost in sending $\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \boldsymbol{\omega}_j)|_{\mathbf{v}_{i,k+1}}$ (which is not present if we send the PDF of $\boldsymbol{\omega}_{i,k}$ at each time k) is completely justifiable.

Remark 2: The need of sending $\nabla_{\mathbf{x}} f_{i,k,j}(\mathbf{x}, \boldsymbol{\omega}_j)|_{\mathbf{v}_{i,k+1}}$ can be properly tuned by selecting small or high ν . In the limit $\nu \rightarrow 0$, there is no need of sending it.

Remark 3: In the policy π_{ij} , each node needs to know Δ_i and Δ_j for all its neighboring nodes. This can be avoided by considering a more conservative policy which substitutes Δ_i and Δ_j with n .

Remark 4: If each of the functions $f_{i,k}(\mathbf{x}, \boldsymbol{\omega})$ is a sum of two functions, one dependent on \mathbf{x} and one linearly dependent on $\boldsymbol{\omega}$, then it is easy to see that the policy π_{ij} becomes checking the utility metric $\mathcal{U}_{S,ij}^{(1)}(\boldsymbol{\omega}_k|\boldsymbol{\omega}')$ only, as we will see in one of the numerical examples.

V. NUMERICAL EVALUATIONS

We now look at two different numerical examples to assess the proposed algorithm in terms of communication overhead and general performance.

A. Least-squares estimation in sensor networks

As a first example, we consider an estimation problem, where each of the computing nodes needs to estimate the state $\mathbf{x} \in [-1/2, 1/2]^d$ of a slowly time-varying process. Each node relies on its own measurements $\mathbf{z}_{i,k} \in \mathbb{R}^d$ that are linked to the state via the measurement equation,

$$\mathbf{z}_{i,k} = \mathbf{H}_i(\boldsymbol{\omega})\mathbf{x}_k + \mathbf{w}_i, \quad (10)$$

where \mathbf{H}_i is a squared stochastic matrix assumed to be full rank for the sake of simplicity. The noise term \mathbf{w}_i is also zero-mean Gaussian and has covariance \mathbf{R}_i . We assume that each of the \mathbf{H}_i 's has the form,

$$\mathbf{H}_i(\boldsymbol{\omega}) = \bar{\mathbf{H}}_i + \mathbf{I}_d \sum_{j \in N_i \cup i} c_j \omega_j, \quad (11)$$

while we assume that each scalar $\omega_i \in \mathbb{R}_+$ is drawn from a truncated Rayleigh distribution with scale parameter $\sigma_{i,k} > 0$ over the compact set $[0, 3]$, slowly varying with time according to a truncated first-order model,

$$\sigma_{i,k+1} = \text{P}_{[0,3]} [\sigma_{i,k} + \rho_i \sin(a_i k + b_i) + r_i], \quad r_i \sim \mathcal{N}(0, P_i).$$

We consider the following time-varying least-squares estimator for the process \mathbf{x} at each time k

$$\underset{\mathbf{x}_k \in [-1/2, 1/2]^d}{\text{minimize}} \quad \mathbb{E}_{\boldsymbol{\omega}, k} \left[\sum_{i=1}^n \|\mathbf{z}_{i,k} - \mathbf{H}_i(\boldsymbol{\omega})\mathbf{x}_k\|_{\mathbf{R}_i^{-1}}^2 \right] \quad (12)$$

and we approximate the expectation operator as a finite sum over N Monte Carlo samples of the stochastic variables $\boldsymbol{\omega}$.

Problem (12) is the optimization program that the computing nodes have to solve in a distributed way using the proposed algorithm, and it is indeed a particular instance of (1). Distributed estimation tasks of the type of (12) may arise in synthetic aperture radar (SAR) array scenarios, for example [31], [32]. In SAR arrays, noise can enter in the measurements in a multiplicative fashion (speckle noise) and it can be modeled with a Rayleigh PDF. A coupled measurement equation with of the form (10)-(11) can be used to model correlated noise within the neighboring sensors. Other similar problems arise when sensors have access to pair-wise relative measurements.

The simulation that we are going to present has the following parameters¹: the dimension of the state vector is $d = 2$, the number of sensors is $n = 15$, the true state evolves via a linear discrete time-invariant process (unknown to the sensor nodes),

$$\mathbf{x}_{k+1}^{\text{true}} = \Phi \mathbf{x}_k^{\text{true}} + \mathbf{n},$$

where Φ is the state transition matrix, while \mathbf{n} is the process noise, assumed zero-mean gaussian with covariance \mathbf{Q} and in particular, $\Phi = [.99, 0.01; 0, 1]$, $\mathbf{Q} = \xi_Q \mathbf{I}_2$, $\bar{\mathbf{H}}_i = \mathbf{I}_2$, $\mathbf{R}_i = \xi_R \mathbf{I}_2$, $\xi_Q = \xi_R = 1\text{e-}6$, $\sigma_{i,0} = \max\{i/n + 0.3(\mathcal{U}_{[0,1]} - 0.3), 0.001\}$, $P_i = 1\text{e-}2$, $a_i = i\pi/200$, $b_i = 200\pi\rho_i$, $\rho_i \sim \mathcal{N}(0, 1)$, and the number of Monte-Carlo samples for the stochastic variable $\boldsymbol{\omega}$ is $N = 5000$. We also consider s_{ij} to be the same for all the links. Also, we scale (12) by multiplying by ξ_R for better numerical stability. For the problem at hand, we can easily compute the bound on β (we set $\beta = 1/n - 1\text{e-}4$), while a suitable value for α has to be derived by trial-and-error as usually happens in distributed optimization (since the constants m_f and L are difficult to obtain in practice). Guided by the stochastic-free case, we set $\alpha = 1/400$ which works well in the simulation scenarios. Furthermore, we can explicitly write the gradient $\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \boldsymbol{\omega}_j)|_{\mathbf{v}_{i,k+1}}$ of the scaled problem as,

$$\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \boldsymbol{\omega}_j)|_{\mathbf{v}_{i,k+1}} = 2 (\mathbf{H}_i(\boldsymbol{\omega})^\top \mathbf{H}_i(\boldsymbol{\omega}) \mathbf{v}_{i,k+1} - \mathbf{z}_{i,k}^\top \mathbf{H}_i(\boldsymbol{\omega}))$$

so that with the policy π_{ij} we need to send either the scalar parameter $\sigma_{i,k}$ or the couple $(\mathbf{v}_{i,k+1}, \mathbf{z}_{i,k})$, thus 4 scalar values. Finally, the optimal solution of (12) for comparison is also computable in closed form.

In this simulation example, we look at loosely connected, highly asynchronous sensor nodes with $s_{ij} = 0.3$, $\gamma = 0.98$, and a high noise dependence with $c_i = 1$ in \mathbf{H}_i for all sensors.

The results show convergence of the proposed scheme to an error floor dependent on ϵ (Figure 2). Furthermore, it is possible to appreciate the decrease in communication overhead when selecting different error levels w.r.t. the every-time communication strategy, where each node sends to all its neighbors

¹The code of the simulation will be made available on-line, so that the interested reader can reproduce the results and change the parameters at will.




Fig. 2. Distributed least-squares: performance in terms of error w.r.t. the optimizer. The results have been averaged over 25 realizations.




Fig. 3. Distributed least-squares: performance in terms of PDF communication. The results have been averaged over 25 realizations.

its updated PDF at every time k (Figure 3). In particular, even selecting $\epsilon = 5, \nu = .25\epsilon$ yields acceptable accuracy with a significant reduction of communication. (We notice that the utility-based communication can have a communication count greater than every-time, since in utility-based the nodes may need to send also the gradient of the local cost functions).

Remark 5: More advanced examples can be derived from this least-squares estimation problem. For instance, moving horizon strategies can be considered in the context of maximum a posteriori estimation, as in [7]. One thing to pay attention to is the possible dependence of the cost function on past (estimated) values of the state, i.e., $y_{i,k-1}$; in this case, although the proposed algorithm still converges (under the given assumptions), the physical meaning of the limit point may be less clear.

B. Waypoint generation in robotic networks

The second example is a networked formation control problem, where a number of mobile nodes needs to track a defined point in space and maintain a certain formation. The example is inspired by [33] and has the added aim to show (i) that the proposed algorithm can work with partially overlapping decision variables \mathbf{x} (i.e., there is no need for each of the computing nodes to agree on the total decision variable \mathbf{x} but only on subsets of it), and (ii) how the policy π_{ij} gets simplified in the case of noise entering linearly in the cost function (see Remark 4).

We consider $n = 16$ mobile nodes that have a fixed connection structure and need to track a squared pattern figure in two dimensions (Figure 4). At a given discrete time k , each mobile node i needs to compute a waypoint $\mathbf{x}_{i,k}$ where to head to, this waypoint depends on the current value of the reference point $\mathbf{x}_{i,k}^{\text{ref}}$ and on the neighboring waypoint/reference values. In addition, we consider each of the reference points to be known to the nodes with some degree of uncertainty, as

$$\mathbf{x}_{i,k}^{\text{ref}} = \bar{\mathbf{x}}_{i,k}^{\text{ref}} + \boldsymbol{\omega}_{i,k},$$

where $\bar{\mathbf{x}}_{i,k}^{\text{ref}}$ is the actual reference and $\boldsymbol{\omega}_{i,k}$ is drawn from a given (estimated) PDF. Putting this together, the computing mobile nodes have to solve the optimization problem

$$\underset{\mathbf{x}_k \in X_k}{\text{minimize}} \mathbb{E}_{\boldsymbol{\omega},k} \left[\sum_{i \in V} \left(\theta \|\mathbf{x}_{i,k} - \mathbf{x}_{i,k}^{\text{ref}}\|^2 + \sum_{j \in N_i} \|\mathbf{x}_{i,k} - \mathbf{x}_{j,k} - (\mathbf{x}_{i,k}^{\text{ref}} - \mathbf{x}_{j,k}^{\text{ref}})\|^2 \right) \right], \quad (13)$$

where \mathbf{x}_k is the stacked version of the all $\mathbf{x}_{i,k}$'s and $\theta > 0$ a chosen scaling factor.

Problem (13) is an instance of (1). Problems like (13) arise in multi-agent scenarios where each mobile node reference point knowledge is affected by stochastic disturbances. This added term can be a pre-compensation (before running the trajectory generation algorithm) of known disturbances acting during the time frame in which the mobile node aims at reaching the waypoint. These disturbances can be internal (e.g., motors), or external (e.g., environment). Environment noise can be for example wind, if the nodes are sufficiently far apart so that the wind variations at the nodes w.r.t. the average value are uncorrelated with each other. Wind is often modeled as a Weibull distribution, and therefore with this in mind, we model each of the stochastic variables $\boldsymbol{\omega}_{i,k}$'s as drawn from a Weibull PDF varying in time (there is no need for the compactness assumption for Ω in this case, as we will see). In particular, $\boldsymbol{\omega}_{i,k} = \mathbf{1}_2 \omega_{i,k}$, $\omega_{i,k} \sim \mathcal{W}(\lambda_{i,k}, \sigma_{i,k})$,

$$\lambda_{i,k} = \lambda_{i,k-1} (1 + \cos(r + \phi_{\lambda,i})/1.3) + \omega_{\max,i},$$

$$\sigma_{i,k} = \sigma_{i,k-1} (1 + \cos(r + \phi_{\sigma,i})/1.3) + \omega_{\max,i},$$

where, $\mathcal{W}(\lambda, \sigma)$ is the Weibull distribution with scalar parameters λ and σ , $\phi_{\lambda,i}$ and $\phi_{\sigma,i}$ are drawn from $\mathcal{U}_{[0,2\pi]}$, and r and $\omega_{\max,i}$ are given scalars.

The reference states ($\bar{\mathbf{x}}_{i,k}^{\text{ref}}$) evolve along circular trajectories with constant angular velocity r . The other parameters of the simulation example are: $s_{ij} = 0.7$, $\theta = 0.5$, $X_k = [-50, 50]^{32}$, $\sigma_{i,0} = 4 + \mathcal{U}_{[0,4]}$, $\lambda_{i,0} = 4/n\omega_{i,\max}(2 + \mathcal{U}_{[0,1]})$. The step

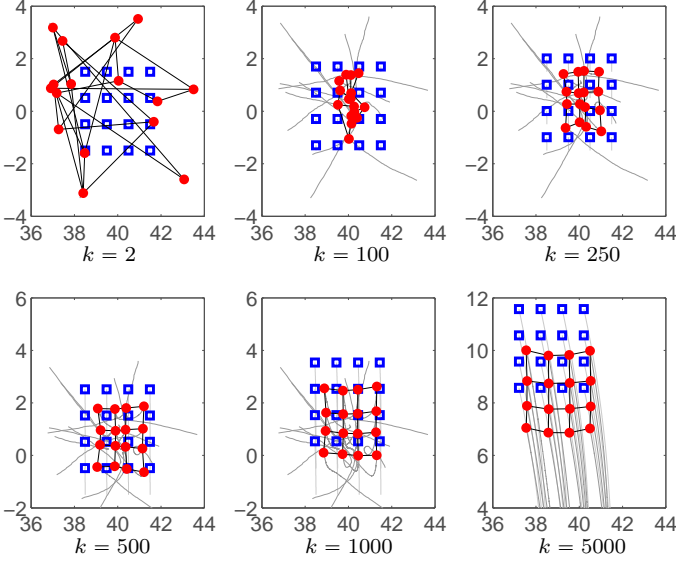


Fig. 4. Snapshots of the algorithm's waypoint generation (red points) and reference ones (blue squares) in the case of no noise.

sizes α and β are determined according to Theorem 1, in fact, in this example, the bounds are analytically computable. For the simulations we set $\alpha = 1.7\theta/(2(\theta + \lambda_n(\bar{\mathbf{W}})/s_{ij}))^2$ and $\beta = 0.13/n$. We can write the gradient $\nabla f_{i,k}$ in closed form as

$$\nabla_{\mathbf{x}} f_{i,k} = 2\theta(\mathbf{x}_{i,k} - \mathbf{x}_{i,k}^{\text{ref}}) + 2 \sum_{j \in N_i} \mathbf{x}_{i,k} - \mathbf{x}_{j,k} - (\mathbf{x}_{i,k}^{\text{ref}} - \mathbf{x}_{j,k}^{\text{ref}}),$$

and since the gradient depends linearly on ω_k (through $\mathbf{x}_{i,k}^{\text{ref}}$), we can simplify the policy π_{ij} by noticing that

$$\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \sum_{j \in N_i} 2\|\hat{\omega}_j' - \hat{\omega}_{j,k}\| \leq \epsilon/2|X_{k,k-1}|, \quad (14)$$

where $\hat{\cdot}$ represents the mean operator; therefore the policy π_{ij} boils down to sending the updated PDF only when condition (14) does not hold.

Example 1. In the first example we consider a simulation scenario without noise. The purpose is to display the performance of the algorithm in its ideal case, when both the gradient is computed exactly and the stochastic program reduces to a deterministic optimization problem (and no compactness assumption on X_k is required). We select the angular velocity as $r = 0.5/40\alpha$, and run the distributed asynchronous time-varying optimization problem up to $k = 5000$. By using snapshots of the agents trajectories we show the algorithm's behavior (Figure 4). The blue squares are the reference waypoints, while the red points are the agent-computed waypoints at discrete time k . The lines represent the possible connection among agents (which are not time-varying [33]). As we further see, the convergence performance is in line with the asymptotical bound of Theorem 1, which is rather tight in this particular case (Figure 5).

Example 2. Stochastic noise is introduced in the second simulation example. We consider the same angular velocity and same α and β . We select $\omega_{i,\max}$ as $2i/n$ and we run Algorithm 1 for different choices of ϵ for the utility-based policy π_{ij} . To simulate a sudden change in the noise (e.g.,




Fig. 5. Convergence performance of the algorithm without noise.




Fig. 6. Distributed waypoint generation: performance in terms of error w.r.t. the optimizer. The results have been averaged over 25 realizations.

a wind gust), we double $\lambda_{i,k}$ in the range $k \in [2500, 3250]$. As we can see in Figures 6 and 7, even considering limited communication exchange, Algorithm 1 with policy π_{ij} performs almost as if all the nodes were communicating their PDF at each time step. We see that the bound of Theorem 1 for $\epsilon = .02$ is tighter than for $\epsilon = .2$, as expected (see Remark 1); nonetheless, the utility-based policy works as designed.

VI. PROOFS

A. Setting Up the Analysis

We define $\mathbf{g}_{i,k}$ as the stochastic gradient of $f_{i,k}$ at $\mathbf{v}_{i,k+1}$,

$$\mathbf{g}_{i,k} = \mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega) | \mathbf{v}_{i,k+1}], \quad (15)$$




Fig. 7. Distributed waypoint generation: performance in terms of PDF communication. The results have been averaged over 25 realizations. (Only one out of every 25 time instances is shown for clarity).

a quantity that will be used to characterize how large “ ϵ ” is for the stochastic ϵ -(sub)gradient $\tilde{\mathbf{g}}_{i,k}$. We also define the stochastic gradient of $f_{i,k}$ at the optimizer of (1), i.e., $\mathbf{x}_k^* = \mathbf{y}_{i,k}^* = \mathbf{v}_{i,k+1}^*$, as,

$$\mathbf{g}_{i,k}^* = \mathbb{E}_{\omega,k}[\nabla_{\mathbf{x}} f_{i,k}(\mathbf{x}, \omega)|_{\mathbf{y}_{i,k}^*}]. \quad (16)$$

Lemma 1: Let $X_{k,k-1}$ be defined as in Theorem 1; let $\mathbf{v}_{i,k+1}$ be defined as (5a), let $\tilde{\mathbf{g}}_{i,k}$, $\mathbf{g}_{i,k}$ and $\mathbf{g}_{i,k}^*$ be defined as (5b), (15), and (16), respectively. Under Assumption 1 and for $\beta < 1/n$, then

- (a) the supporting variable $\mathbf{v}_{i,k+1} \in X_{k-1}$;
- (b) if $\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \epsilon/2|X_{k,k-1}|$ holds, then $\tilde{\mathbf{g}}_{i,k}$ is a stochastic ϵ -(sub)gradient of $f_{i,k}(\mathbf{x}, \omega)$ at $\mathbf{v}_{i,k+1}$;
- (c) at each discrete time step k , the inequality $\|\mathbf{g}_{i,k} - \mathbf{g}_{i,k}^*\| \leq L\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|$ holds.

Proof. (a) The variable $\mathbf{v}_{i,k+1}$ is generated in (5a) as

$$\mathbf{v}_{i,k+1} = \left(1 - \beta \sum_{j=1, j \neq i}^n [\mathbf{A}_k]_{i,j}\right) \mathbf{y}_{i,k} + \beta \sum_{j=1, j \neq i}^n [\mathbf{A}_k]_{i,j} \mathbf{y}_{j,k} = \sum_{j=1}^n \theta_j \mathbf{y}_{j,k},$$

with $\sum_{j=1}^n \theta_j = 1$, $\theta_j, j \neq i \geq 0$. Furthermore,

$$\sum_{j=1, j \neq i}^n [\mathbf{A}_k]_{i,j} < n < 1/\beta,$$

and thus it is also true that $\theta_i = 1 - \beta \sum_{j=1, j \neq i}^n [\mathbf{A}_k]_{i,j} > 0$. This means that $\mathbf{v}_{i,k+1}$ is a convex combination of $\mathbf{y}_{i,k}$ and since $\mathbf{y}_{i,k} \in X_{k-1}$, then $\mathbf{v}_{i,k+1} \in X_{k-1}$.

(b) By the definition of the ϵ -(sub)gradient given in (2) and by the definition of the gradient, by the fact that the domain of the $f_{i,k}$'s needs to be at least $X_{k,k-1} \times \Omega$ by assumption, it has to be

$$(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k})^\top (\mathbf{v} - \mathbf{v}_{i,k+1}) + \mathbf{g}_{i,k}^\top (\mathbf{v} - \mathbf{v}_{i,k+1}) \leq \mathbb{E}_{\omega,k}[f_{i,k}(\mathbf{v}, \omega) - f_{i,k}(\mathbf{v}_{i,k+1}, \omega)] + \epsilon, \quad (17)$$

for all $\mathbf{v} \in X_{k,k-1}$, $\epsilon \geq 0$. Since by definition, the gradient verifies the same condition of (2) with $\epsilon = 0$, a sufficient condition for (17) is

$$\|(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k})^\top (\mathbf{v} - \mathbf{v}_{i,k+1})\| \leq \epsilon,$$

and yet another sufficient condition for this latter inequality to hold is

$$\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \frac{\epsilon}{2|X_{k,k-1}|}.$$

(c) This claimed inequality comes from the Lipschitz assumption applied to the gradients $\mathbf{g}_{i,k}$ and $\mathbf{g}_{i,k}^*$ of the points $\mathbf{v}_{i,k+1} \in X_{k-1}$ and $\mathbf{y}_{i,k}^* \in X_{k-1}$. ■

Lemma 2: With the same definitions of Theorem 1 and under Assumption 1, for any $\mu > 0$, if $\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| \leq \epsilon/2|X_{k,k-1}|$ for all $i, k \geq 1$, then the sequence $\{\mathbf{y}_{i,k}\}$ generated by the proposed algorithm is bounded as

$$\|\mathbf{y}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 \leq (1 + \mu)(1 + \alpha^2 L^2 - \alpha m_f) \|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 + (1 + \mu)\alpha\epsilon \left(\alpha \frac{\epsilon}{4|X_{k,k-1}|^2} + 2\alpha L + 2 \right) + (1 + 1/\mu)\alpha^2 G^2.$$

Proof. By the definition of the update rule (5c) and the optimal gradient $\mathbf{g}_{i,k}^*$ given in (16), we can write

$$\begin{aligned} \|\mathbf{y}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 &= \|\mathbf{P}_{X_k}[\mathbf{v}_{i,k+1} - \alpha \tilde{\mathbf{g}}_{i,k}] - \mathbf{P}_{X_k}[\mathbf{y}_{i,k} - \alpha \sum_{j \in V} \mathbf{g}_{j,k}^*]\|^2 \\ &\leq \left\| \mathbf{v}_{i,k+1} - \alpha(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*) - \mathbf{y}_{i,k}^* + \alpha \sum_{j \in V, j \neq i} \mathbf{g}_{j,k}^* \right\|^2, \end{aligned} \quad (18)$$

where we have used $\mathbf{y}_{i,k}^* = \mathbf{P}_{X_k}[\mathbf{y}_{i,k}^* - \alpha \sum_{j \in V} \mathbf{g}_{j,k}^*]$, and the non-expansivity property of the projection. If we use the Peter-Paul inequality² and the boundedness of the optimal gradient (Eq. (6)), we obtain

$$\|\mathbf{y}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 \leq (1 + \mu) \|\mathbf{v}_{i,k+1} - \alpha(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*) - \mathbf{y}_{i,k}^*\|^2 + (1 + 1/\mu)\alpha^2 G^2, \quad (19)$$

If we expand the first part of the bound in (19), we obtain

$$\begin{aligned} \|\mathbf{v}_{i,k+1} - \alpha(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*) - \mathbf{y}_{i,k}^*\|^2 &= \|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 + \\ &\alpha^2 \|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*\|^2 - 2\alpha(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*)^\top (\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*). \end{aligned} \quad (20)$$

By using Lemma 1, Assumption 1, and the definition of ϵ -(sub)gradient, we can now bound the following terms,

$$\begin{aligned} \|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*\|^2 &\leq (\|\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}\| + \|\mathbf{g}_{i,k} - \mathbf{g}_{i,k}^*\|)^2 \leq \\ &\left(\frac{\epsilon}{2|X_{k,k-1}|} + L\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\| \right)^2 \leq \\ &\frac{\epsilon^2}{4|X_{k,k-1}|^2} + 2L\epsilon + L^2\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 \end{aligned}$$

where we bounded $\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\| \leq 2|X_{k,k-1}|$, and by strong convexity (recall $\mathbf{v}_{i,k+1} \in X_{k-1}$)

$$\begin{aligned} -(\tilde{\mathbf{g}}_{i,k} - \mathbf{g}_{i,k}^*)^\top (\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*) &= \\ \tilde{\mathbf{g}}_{i,k}^\top (\mathbf{y}_{i,k}^* - \mathbf{v}_{i,k+1}) + \mathbf{g}_{i,k}^{*\top} (\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*) &\leq \\ \mathbb{E}_{\omega,k}[f_{i,k}(\mathbf{y}_{i,k}^*, \omega) - f_{i,k}(\mathbf{v}_{i,k+1}, \omega)] + \mathbf{g}_{i,k}^{*\top} (\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*) + \epsilon & \\ \leq -\frac{m_f}{2}\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 + \epsilon. \end{aligned}$$

Combining these results with (20) and (19) we obtain,

$$\begin{aligned} \|\mathbf{y}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 &\leq (1 + \mu) \left[\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 + \right. \\ &\alpha^2 \left(\frac{\epsilon^2}{4|X_{k,k-1}|^2} + 2L\epsilon + L^2\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 \right) + \\ &\left. \alpha \left(-m_f\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2 + 2\epsilon \right) \right] + (1 + 1/\mu)\alpha^2 G^2, \end{aligned}$$

and therefore the claim is proven. ■

We turn now our attention to the term $\|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2$, which can be bounded as follows.

Lemma 3: Let \mathbf{y}_k be the stacked version of the local decision variables $\mathbf{y}_{i,k}$, and let $\bar{\mathbf{W}} = \mathbb{E}[\mathbf{W}_k]$. Under the same assumptions and with the same definitions of Lemma 2 we have

$$\mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_k^*\|^2] \leq \varphi(\alpha, \bar{\mathbf{W}}, \mu) \mathbb{E}[\|\mathbf{y}_k - \mathbf{y}_k^*\|^2] + \alpha\psi(\alpha, \epsilon, \mu),$$

²I.e., $(a + b)^2 \leq (1 + \mu)a^2 + (1 + 1/\mu)b^2$, for all $a, b \in \mathbb{R}$ and $\mu > 0$.

where

$$\begin{aligned}\varphi(\alpha, \bar{\mathbf{W}}, \mu) &= (1 + \mu)(1 + \alpha^2 L^2 - \alpha m_f)(1 - \beta \lambda_2(\bar{\mathbf{W}})) \\ \psi(\alpha, \epsilon, \mu) &= (1 + \mu)n\epsilon \left(\alpha \frac{\epsilon}{4|X_{k,k-1}|^2} + 2\alpha L + 2 \right) + \\ &\quad (1 + 1/\mu)\alpha n G^2.\end{aligned}$$

Proof. By the definition of $\mathbf{v}_{i,k+1}$ in (5a), we can write its stacked version for all i 's as

$$\mathbf{v}_{k+1} = \mathbf{y}_k - \beta(\mathbf{W}_k \otimes \mathbf{I}_d)\mathbf{y}_k,$$

and since $(\mathbf{W}_k \otimes \mathbf{I}_d)\mathbf{y}_k^* = \mathbf{0}_{nd}$, we can expand

$$\|\mathbf{v}_{k+1} - \mathbf{y}_k^*\|^2 = (\mathbf{y}_k - \mathbf{y}_k^*)^\top (\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d)^2 (\mathbf{y}_k - \mathbf{y}_k^*).$$

We use now the fact that $\beta < 1/n$, which also means³ $\mathbf{I}_{nd} \geq \mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d > 0$, therefore $\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d \geq (\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d)^2$, and thus

$$\begin{aligned}\|\mathbf{v}_{k+1} - \mathbf{y}_k^*\|^2 &= (\mathbf{y}_k - \mathbf{y}_k^*)^\top (\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d)^2 (\mathbf{y}_k - \mathbf{y}_k^*) \leq \\ &(\mathbf{y}_k - \mathbf{y}_k^*)^\top (\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d) (\mathbf{y}_k - \mathbf{y}_k^*).\end{aligned}$$

We take now the expectation of the previous expression

$$\begin{aligned}\mathbb{E}[\|\mathbf{v}_{k+1} - \mathbf{y}_k^*\|^2] &\leq \\ &\mathbb{E}[(\mathbf{y}_k - \mathbf{y}_k^*)^\top (\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d) (\mathbf{y}_k - \mathbf{y}_k^*)],\end{aligned}$$

and since \mathbf{y}_k and \mathbf{W}_k are independent,

$$\begin{aligned}\mathbb{E}[\|\mathbf{v}_{k+1} - \mathbf{y}_k^*\|^2] &\leq \\ &\text{tr}((\mathbf{I}_{nd} - \beta \bar{\mathbf{W}} \otimes \mathbf{I}_d) \mathbb{E}[(\mathbf{y}_k - \mathbf{y}_k^*)(\mathbf{y}_k - \mathbf{y}_k^*)^\top]),\end{aligned}$$

where $\text{tr}(\cdot)$ is the trace operator. And finally, given that $\bar{\mathbf{W}}(\mathbf{y}_k - \mathbf{y}_k^*)$ will be zero only when we reach consensus (i.e., at optimality),

$$\begin{aligned}\text{tr}((\mathbf{I}_{nd} - \beta \bar{\mathbf{W}} \otimes \mathbf{I}_d) \mathbb{E}[(\mathbf{y}_k - \mathbf{y}_k^*)(\mathbf{y}_k - \mathbf{y}_k^*)^\top]) &\leq \\ &(1 - \beta \lambda_2(\bar{\mathbf{W}})) \mathbb{E}[\|\mathbf{y}_k - \mathbf{y}_k^*\|^2],\end{aligned}$$

and the claim follows by combining this with Lemma 2:

$$\begin{aligned}\mathbb{E}\left[\sum_{i \in V} \|\mathbf{y}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2\right] &\leq \\ &\mathbb{E}\left[\sum_{i \in V} (1 + \mu)(1 + \alpha^2 L^2 - \alpha m_f) \|\mathbf{v}_{i,k+1} - \mathbf{y}_{i,k}^*\|^2\right] + \\ &\alpha \psi(\alpha, \epsilon, \mu) \leq \varphi(\alpha, \bar{\mathbf{W}}, \mu) \mathbb{E}[\|\mathbf{y}_k - \mathbf{y}_k^*\|^2] + \alpha \psi(\alpha, \epsilon, \mu).\end{aligned}$$

■

B. Proof of Theorem 1

We are now ready to prove Theorem 1.

Proof. We use again the Peter-Paul inequality, Assumption 3, and Lemma 3 to establish

$$\begin{aligned}\mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_{k+1}^*\|^2] &\leq (1 + \mu) \mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_k^*\|^2] + \\ &(1 + 1/\mu) \mathbb{E}[\|\mathbf{y}_{k+1}^* - \mathbf{y}_k^*\|^2] \leq n \delta_x^2 (1 + 1/\mu) + \\ &(\varphi(\alpha, \bar{\mathbf{W}}, \mu) \mathbb{E}[\|\mathbf{y}_k - \mathbf{y}_k^*\|^2] + \alpha \psi(\alpha, \epsilon, \mu)) (1 + \mu),\end{aligned}$$

³In fact, $\lambda_1(\mathbf{I}_{nd} - \beta \mathbf{W}_k \otimes \mathbf{I}_d) = 1 - \beta \lambda_n(\mathbf{W}_k)$; bounding $\lambda_n(\mathbf{W}_k) \leq n$ from [34], the result follows.

which is true for any scalar $\mu > 0$. In order to ensure convergence of the sequence $\{\mathbf{y}_k\}$, it has to be

$$r := \varphi(\alpha, \bar{\mathbf{W}}, \mu)(1 + \mu) < 1,$$

so that,

$$\begin{aligned}\mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_{k+1}^*\|^2] &\leq r^k \mathbb{E}[\|\mathbf{y}_1 - \mathbf{y}_1^*\|^2] + \\ &\sum_{\tau=1}^k r^\tau (\alpha \psi(\alpha, \epsilon, \mu)(1 + \mu) + n \delta_x^2 (1 + 1/\mu)),\end{aligned}$$

and finally, $\liminf_{k \rightarrow \infty} \mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_{k+1}^*\|^2] =$

$$\begin{aligned}\liminf_{k \rightarrow \infty} \sum_{\tau=1}^k r^\tau (\alpha \psi(\alpha, \epsilon, \mu)(1 + \mu) + n \delta_x^2 (1 + 1/\mu)) \\ \leq \frac{1}{1-r} (\alpha \psi(\alpha, \epsilon, \mu)(1 + \mu) + n \delta_x^2 (1 + 1/\mu)).\end{aligned}$$

To ensure $r < 1$, we need to have

$$(1 + \alpha^2 L^2 - \alpha m_f)(1 - \beta \lambda_2(\bar{\mathbf{W}}))(1 + \mu)^2 < 1.$$

We call $\gamma = 1 - \beta \lambda_2(\bar{\mathbf{W}})$, and by the definition of β and Assumption 4, it is $0 < \gamma < 1$. By choosing $\mu = 1/\sqrt{\gamma} - 1$, we obtain the condition $\varrho := 1 + \alpha^2 L^2 - \alpha m_f < 1$, that is $\alpha < m_f/L^2$, and the limit result

$$\begin{aligned}\liminf_{k \rightarrow \infty} \mathbb{E}[\|\mathbf{y}_{k+1} - \mathbf{y}_{k+1}^*\|^2] &\leq \\ &\frac{1}{1-\varrho} \left(\psi(\alpha, \epsilon, \mu) \frac{\alpha}{\sqrt{\gamma}} + n \delta_x^2 \frac{1}{1-\sqrt{\gamma}} \right),\end{aligned}$$

with $\mu = 1/\sqrt{\gamma} - 1$, and the theorem is proven. ■

C. Proof of Theorem 2

Proof. We start expanding the condition (7) by using the triangle inequality as

$$\begin{aligned}\|\mathbb{E}_{\omega_j, k}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \mathbb{E}_{\omega'_j}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}]\| \\ \leq \|\mathbb{E}_{\omega_j, k}[\nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}] - \mathbb{E}_{\omega'_j}[\nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\| + \\ \|\mathbb{E}_{\omega_j, k}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\| \\ - \mathbb{E}_{\omega'_j}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\|.\end{aligned}$$

The right-most term can be bounded as

$$\begin{aligned}\|\mathbb{E}_{\omega_j, k}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\| \\ - \mathbb{E}_{\omega'_j}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}]\| \leq \\ \int_{\Omega_j} \|\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}} - \nabla \mathbf{x} f'_{i,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}'_i}\| \cdot \\ |p_{\omega_j, k}(\omega) - p_{\omega'_j}(\omega)| d\omega \leq \nu \mathcal{U}_{R,ij}(\nabla \mathbf{x} f_{i,k,j} | \nabla \mathbf{x} f'_{i,j}),\end{aligned}$$

where we have used the policy π_{ij} to bound $|p_{\omega_j, k}(\omega) - p_{\omega'_j}(\omega)|$. Therefore we can write

$$\begin{aligned}\|\mathbb{E}_{\omega_j, k}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}] - \mathbb{E}_{\omega'_j}[\nabla \mathbf{x} f_{i,k,j}(\mathbf{x}, \omega_j)|_{\mathbf{v}_{i,k+1}}]\| \\ \leq \mathcal{U}_{S,ji}^{(1)}(\omega_k | \omega') + \nu \mathcal{U}_{R,ij}(\nabla \mathbf{x} f_{i,k,j} | \nabla \mathbf{x} f'_{i,j}).\end{aligned}$$

And finally, a sufficient condition for (7) is

$$\mathcal{U}_{S,ji}^{(1)}(\omega_k|\omega') \leq (1-\eta) \frac{\epsilon}{2|X_{k,k-1}|\Delta_i},$$

$$\mathcal{U}_{R,ij}(\nabla_x f_{i,k,j}|\nabla_x f'_{i,j}) \leq \frac{\eta}{\nu} \frac{\epsilon}{2|X_{k,k-1}|\Delta_i}$$

for any $\eta \in [0, 1]$; this together with the bound on $\mathcal{U}_{S,j}^{(2)}(\omega_k|\omega')$ yields the claim. ■

VII. CONCLUSIONS

We have proposed a distributed stochastic ϵ -gradient asynchronous algorithm to optimize a rather general convex separable time-varying stochastic program. To alleviate the possibly high communication demands among the nodes to track variations in the PDF's of the stochastic variables, we have devised a utility-based policy to trigger the said communication. The overall scheme converges linearly to an error bound whose size depends on the constant stepsize α , on ϵ , and the variability in time of the optimizer. Initial simulation results are encouraging and well display the added value of the proposed approach.

Many relevant questions require further studies. First of all, the Peter-Paul inequality is known to be a loose bound; by substituting this inequality with tighter bounds, we expect more accurate asymptotic results. The algorithm should be further modified to allow for less restrictive assumptions, both on the cost functions and on the stochastic variables. Finally, in order to be able to characterize dynamical constraints, it will be interesting to explicitly include in the analysis cost functions that depend on past variables generated locally, e.g., $f_{i,k}$ should also depend on $y_{i,k-1}$, as pointed out in Remark 5.

REFERENCES

- [1] L. Xiao, "Dual Averaging Methods for Regularized Stochastic Learning and Online Optimization," *Journal of Machine Learning Research*, vol. 11, pp. 2543 – 2596, 2010.
- [2] J. Duchi, M. Jordan, and H. Brendan McMahan, "Estimation, Optimization, and Parallelism when Data is Sparse," in *Proceedings of Neural Information Processing Systems*, Lake Tahoe, US, December 2013, pp. 1 – 18.
- [3] E. van Foecken, L. Kester, and M. van Iersel, "Real-time Common Awareness in Communication Constrained Sensor Systems," in *Proceedings of 12th International Conference on Information Fusion*, Seattle, US, July 2009, pp. 118 – 125.
- [4] M. Ditzel, L. Kester, S. van den Broek, and M. van Rijn, "Cross-layer Utility-based System Optimization," in *Proceedings of 16th International Conference on Information Fusion*, Istanbul, Turkey, July 2013, pp. 507 – 514.
- [5] J. Ueda, L. Odhner, and H. H. Asada, "Broadcast Feedback of Large-Scale, Distributed Stochastic Control Systems Inspired by Biological Muscle Control," in *Proceedings of the American Control Conference*, New York City, US, July 2007, pp. 1317 – 1322.
- [6] S. Boyd, M. Mueller, B. O'Donoghue, and Y. Wang, "Performance Bounds and Suboptimal Policies for Multi-Period Investment," *Foundations and Trends® in Optimization*, vol. 1, no. 1, pp. 1 – 72, 2014.
- [7] F. Y. Jakubiec and A. Ribeiro, "D-MAP: Distributed Maximum a Posteriori Probability Estimation of Dynamic Systems," *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 450 – 466, 2013.
- [8] C. Zhao, U. Topcu, and S. H. Low, "Fast Load Control with Stochastic Frequency Measurement," in *Proceedings of the IEEE Power and Energy Society General Meeting*, San Diego, US, July 2012, pp. 1 – 8.
- [9] S. H. Low, "Convex Relaxation of Optimal Power Flow, Part I: Formulations and Equivalence," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 15 – 27, 2014.
- [10] C. Nowakowski and S. E. Shladover, D. Cody, et al., "Cooperative Adaptive Cruise Control: Testing Drivers' Choices of Following Distances," University of California, Tech. Rep., 2011, California PATH Research Report UCB-ITS-PRR-2011-01, <http://www.path.berkeley.edu>.
- [11] S.-Y. Tu and A. H. Sayed, "Mobile Adaptive Networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 649 – 664, 2011.
- [12] A. Simonetto, T. Keviczky, and R. Babuska, "Constrained Distributed Algebraic Connectivity Maximization in Robotic Networks," *Automatica*, vol. 49, no. 5, pp. 1348 – 1357, 2013.
- [13] K. Srivastava and A. Nedić, "Distributed Asynchronous Constrained Stochastic Optimization," *IEEE Transactions on Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772 – 790, 2011.
- [14] J. C. Duchi, A. Agarwal, and M. Wainwright, "Dual Averaging for Distributed Optimization: Convergence Analysis and Network Scaling," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 592 – 606, 2012.
- [15] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1 – 122, 2011.
- [16] M. Kamgarpour and C. Tomlin, "Convergence Properties of a Decentralized Kalman Filter," in *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, December 2008, pp. 3205 – 3210.
- [17] P. Braca, S. Marano, V. Matta, and P. Willett, "Asymptotic Optimality of Running Consensus in Testing Binary Hypotheses," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 814 – 825, 2010.
- [18] F. S. Cattivelli and A. H. Sayed, "Diffusion Strategies for Distributed Kalman Filtering and Smoothing," *IEEE Transactions on Automatic Control*, vol. 55, no. 9, pp. 2069 – 2084, 2010.
- [19] M. Farina, G. Ferrari-Trecate, and R. Scattolini, "Distributed Moving Horizon Estimation for Linear Constrained Systems," *IEEE Transactions on Automatic Control*, vol. 55, no. 11, pp. 2462 – 2475, 2010.
- [20] D. Bajovic, D. Jakovetic, J. Xavier, B. Sinopoli, and J. M. F. Moura, "Distributed Detection via Gaussian Running Consensus: Large Deviations Asymptotic Analysis," *IEEE Transactions on Signal Processing*, vol. 59, no. 9, pp. 4381 – 4396, 2011.
- [21] M. M. Zavlanos, A. Ribeiro, and G. J. Pappas, "Network Integrity in Mobile Robotic Networks," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 3 – 18, 2013.
- [22] R. L. G. Cavalcante and S. Stanczak, "A Distributed Subgradient Method for Dynamic Convex Optimization Problems Under Noisy Information Exchange," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 243 – 256, 2013.
- [23] Q. Ling and A. Ribeiro, "Decentralized Dynamic Optimization Through the Alternating Direction Method of Multipliers," *IEEE Transactions on Signal Processing*, vol. 62, no. 5, pp. 1185 – 1197, 2014.
- [24] P. Wan and M. D. Lemmon, "Event-Triggered Distributed Optimization in Sensor Networks," in *Proceedings of the International Conference on Information Processing in Sensor Networks*, San Francisco, US, April 2009, pp. 49 – 60.
- [25] M. Zhong and C. G. Cassandras, "Asynchronous Distributed Optimization With Event-Driven Communication," *IEEE Transactions on Automatic Control*, vol. 55, no. 12, pp. 2735 – 2750, 2010.
- [26] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, "An introduction to event-triggered and self-triggered control," in *Proceedings of the IEEE 51st Conference on Decision and Control*, Maui, US, December 2012, pp. 3270 – 3285.
- [27] S. Trimpe and R. D'Andrea, "Event-based State Estimation with Variance-based Triggering," in *Proceedings of the IEEE 51st Conference on Decision and Control*, Maui, USA, December 2012, pp. 6583 – 6590.
- [28] R. Arroyo-Valles, A. G. Marques, and J. Cid-Sueiro, "Optimal Selective Transmission under Energy Constraints in Sensor Networks," *IEEE Transactions on Mobile Computing*, vol. 8, no. 11, pp. 1524 – 1538, 2009.
- [29] A. Razavi, W. Zhang, and Z.-Q. Luo, "Distributed Optimization in an Energy-Constrained Network: Analog Versus Digital Communication Schemes," *IEEE Transactions on Information Theory*, vol. 59, no. 3, pp. 1803 – 1817, 2013.
- [30] B. T. Polyak, *Introduction to Optimization*. Optimization Software, Inc., 1987.
- [31] M. Pardini, F. Lombardini, and F. Gini, "The Hybrid Cramér Rao Bound on Broadside DOA Estimation of Extended Sources in Presence of Array Errors," *IEEE Transactions on Signal Processing*, vol. 56, no. 4, pp. 1726 – 1730, 2008.
- [32] J. M. Bioucas-Dias and M. A. T. Figueiredo, "Multiplicative Noise Removal Using Variable Splitting and Constrained Optimization," *IEEE Transactions on Image Processing*, vol. 19, no. 7, pp. 1720 – 1730, 2010.
- [33] F. Borrelli and T. Keviczky, "Distributed LQR Design for Identical Dynamically Decoupled Systems," *IEEE Transaction on Automatic Control*, vol. 53, no. 8, pp. 1901 – 1912, 2008.
- [34] O. Rojo, "A nontrivial upper bound on the largest laplacian eigenvalue of weighted graphs," *Linear Algebra and Its Applications*, vol. 420, no. 2 – 3, pp. 625 – 633, 2007.